



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :
G06F 15/16, 11/00, 13/00, 15/76

A1

(11) International Publication Number: WO 95/15529

(43) International Publication Date: 8 June 1995 (08.06.95)

(21) International Application Number: PCT/US94/13350

(22) International Filing Date: 15 November 1994 (15.11.94)

(30) Priority Data:
08/159,783 1 December 1993 (01.12.93) US

(71) Applicant: MARATHON TECHNOLOGIES CORPORATION [US/US]; 70 Birch Hill Road, Stow, MA 01775 (US).

(72) Inventors: BISSETT, Thomas, D.; 78 Sunset Drive, Northborough, MA 01532 (US). FIORENTINO, Richard, D.; 43 Aaron Way, Carlisle, MA 01741 (US). GLORIOSO, Robert, M.; 70 Birch Hill Road, Stow, MA 01775 (US). MCCAULEY, Diane, T.; 190 Elm Street, Hopkinton, MA 01748 (US). MCCOLLUM, James, D.; 275 Swift Road, Whitinsville, MA 01588 (US). TREMBLAY, Glenn, A.; 139 South Street, Upton, MA 01568 (US). TROLANI, Mario; 57 Walnut Street, Newton, MA 02160 (US).

(74) Agent: WALPERT, Gary, A.; Fish & Richardson, 225 Franklin Street, Boston, MA 02110-2804 (US).

(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published


With international search report.

Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: FAULT RESILIENT/FAULT TOLERANT COMPUTING

(57) Abstract

A method of synchronizing at least two computing elements (CE1, CE2) that each have clocks that operate asynchronously of the clocks of the other computing elements includes selecting one or more signals, designated as meta time signals, from a set of signals produced by the computing elements (CE1, CE2), monitoring the computing elements (CE1, CE2) to detect the production of a selected signal by one of the computing elements (CE1), waiting for the other computing elements (CE2) to produce a selected signal, transmitting equally valued time updates to each of the computing elements, and updating the clocks of the computing elements (CE1, CE2) based on the time updates. In a second aspect of the invention, fault resilient, or tolerant, computers (200) are produced by designating a first processor as a computing element (204), designating a second processor (202) as a controller, connecting the computing element (204) and the controller (202) to produce a modular pair, and connecting at least two modular pairs to produce a fault resilient or fault tolerant computer (200). Each computing element (202, 204) of the computer (200) performs all instructions in the same number of cycles as the other computing elements (202, 204). The computer systems include one or more controllers (202) and at least two computing elements (204).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LJ	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

- 1 -

FAULT RESILIENT/FAULT TOLERANT COMPUTINGBackground of the Invention

5 The invention relates to fault resilient and fault tolerant computing methods and apparatus.

 Fault resilient computer systems can continue to function in the presence of hardware failures. These systems operate in either an availability mode or an integrity mode, but not both. A system is "available" when a hardware failure does not cause unacceptable delays in user access, and a system operating in an availability mode is configured to remain online, if possible, when faced with a hardware error. A system has data integrity when a hardware failure causes no data loss or corruption, and a system operating in an integrity mode is configured to avoid data loss or corruption, even if it must go offline to do so.

 Fault tolerant systems stress both availability and integrity. A fault tolerant system remains available and retains data integrity when faced with a single hardware failure, and, under some circumstances, with multiple hardware failures.

 Disaster tolerant systems go one step beyond fault tolerant systems and require that loss of a computing site due to a natural or man-made disaster will not interrupt system availability or corrupt or lose data.

 Prior approaches to fault tolerance include software checkpoint/restart, triple modular redundancy, and pair and spare.

 Checkpoint/restart systems employ two or more computing elements that operate asynchronously and may execute different applications. Each application periodically stores an image of the state of the computing element on which it is running (a checkpoint).

- 2 -

When a fault in a computing element is detected, the checkpoint is used to restart the application on another computing element (or on the same computing element once the fault is corrected). To implement a

5 checkpoint/restart system, each of the applications and/or the operating system to be run on the system must be modified to periodically store the image of the system. In addition, the system must be capable of "backtracking" (that is, undoing the effects of any

10 operations that occurred subsequent to a checkpoint that is being restarted).

With triple modular redundancy, three computing elements run the same application and are operated in cycle-by-cycle lockstep. All of the computing elements

15 are connected to a block of voting logic that compares the outputs (that is, the memory interfaces) of the three computing elements and, if all of the outputs are the same, continues with normal operation. If one of the outputs is different, the voting logic shuts down the

20 computing element that has produced the differing output. The voting logic, which is located between the computing elements and memory, has a significant impact on system speed.

Pair and spare systems include two or more pairs

25 of computing elements that run the same application and are operated in cycle-by-cycle lockstep. A controller monitors the outputs (that is, the memory interfaces) of each computing element in a pair. If the outputs differ, both computing elements in the pair are shut down.

30 Summary of the Invention

According to the invention, a fault resilient and/or fault tolerant system is obtained through use of at least two computing elements ("CEs") that operate asynchronously in real time (that is, from cycle to

35 cycle) and synchronously in so-called "meta time." The

- 3 -

CEs are synchronized at meta times that occur often enough so that the applications running on the CEs do not diverge, but are allowed to run asynchronously between the meta times. For example, the CEs could be

5 synchronized once each second and otherwise run asynchronously. Because the CEs are resynchronized at each meta time, the CEs are said to be operating in meta time lockstep.

In particular embodiments, meta times are defined

10 as the times at which the CEs request I/O operations. In these embodiments, the CEs are synchronized after each I/O operation and run asynchronously between I/O operations. This approach is applicable to systems in which at least two asynchronous computing elements

15 running identical applications always generate I/O requests in the same order. This approach can be further limited to resynchronization after only those I/O requests that modify the processing environment (that is, write requests).

20 Meta time synchronization according to the invention is achieved through use of a paired modular redundant architecture that is transparent to applications and operating system software. According to this architecture, each CE is paired with a controller,

25 otherwise known as an I/O processor ("IOP"). The IOPs perform any I/O operations requested by or directed to the CEs, detect hardware faults, and synchronize the CEs with each other after each I/O operation. In systems in which I/O requests are not issued with sufficient

30 frequency, the IOPs periodically synchronize the CEs in response to so-called "quantum interrupts" generated by inter-processor interconnect (IPI) modules connected to the CEs.

In another particular embodiment of the invention,

35 rather than synchronizing the CEs based on each

- 4 -

particular I/O operation, the CEs are synchronized based on a window of I/O operations. In this approach, a list of I/O operations is maintained for each CE and the CEs are synchronized whenever a common entry appears in all
5 of the lists. This approach allows flexibility as to the order in which I/O requests are generated.

In yet another exemplary embodiment of the invention, the CEs are synchronized based either on signals that are periodically generated by the operating
10 system or on hardware generated interrupts. For example, in the hardware interrupt approach, a processor of each CE is modified to generate an interrupt every N cycles and the CEs are synchronized in response to those interrupts.

15 Primary components of a paired modular redundant system include software, off-the-shelf IOPs, off-the-shelf CEs, and pairs of customized IPI modules that plug into expansion slots of the IOP and the CE and are interconnected by a cable. Redundant I/O devices can be
20 connected to one or more of the CEs or IOPs to provide redundant I/O and offer features such as volume shadowing of key mass storage devices. A paired modular redundant system can accommodate any I/O device that is compatible with a processor used in implementing an IOP of the
25 system.

The paired modular redundant architecture uses minimal custom software and hardware to enable at least two off-the-shelf computing elements to be combined into a fault resilient or tolerant system that runs industry
30 standard operating systems, such as Windows NT, DOS, OS/2, or UNIX, and unmodified applications. Thus, the architecture can avoid both the high costs and inflexibility of the proprietary operating systems, applications, and processor designs used in the prior
35 art.

- 5 -

Another advantage of the paired modular redundant architecture of the present invention is that it offers a certain degree of software fault tolerance. The majority of software errors are not algorithmic. Instead, most errors are caused by asynchrony between the computing element and I/O devices that results in I/O race conditions. By decoupling I/O requests from the computing elements, the paired modular redundant architecture should substantially reduce the number of so-called "Heisenbug" software errors that result from such asynchrony.

In one aspect, generally, the invention features forming a fault tolerant or fault resilient computer by using at least one controller to synchronize at least two computing elements that each have clocks operating asynchronously of the clocks of the other computing elements. One or more signals, designated as meta time signals, are selected from a set of signals produced by the computing elements. Thereafter, the computing elements are monitored to detect the production of selected signals by one of the computing elements. Once a selected signal is detected, the system waits for the production of selected signals by the other computing elements, and, upon receiving the selected signals, transmits equal time updates to each of the computing elements. The clocks of the computing elements are then updated based on the time updates.

Preferred embodiments of the invention include the features listed below. First, I/O requests are the selected signals. The I/O requests are processed to produce I/O responses that are transmitted with the time updates. In addition to, or instead of, I/O requests, quantum interrupts can be the selected signals. The computing elements count either executed instructions or the cycles of a clock such as the system clock, bus

- 6 -

clock, or I/O clock, and generate quantum interrupts whenever a predefined number of instructions or cycles occurs. When both I/O requests and quantum interrupts are used as the selected signals, the computing elements
5 count the number of instructions or cycles that occur without an I/O request. For example, a computing element could be programmed to generate a quantum interrupt whenever it processes for one hundred cycles without generating an I/O request.

10 In one embodiment, instructions are counted by loading a counter with a predetermined value, enabling the counter with an I/O request, decrementing the value of the counter, and signalling a quantum interrupt when the value of the counter reaches zero. In another
15 approach, debugging features of the processor are used to generate the quantum interrupts.

For fault detection, the selected signals and accompanying data, if any, from each of the computing elements are compared. If they do not match, a signal is
20 generated to indicate that a fault has occurred.

In some embodiments, the computing elements wait for time updates by pausing operation after producing the selected signals. The computing elements resume operation upon receipt of the time updates. In other
25 embodiments, the computing elements continue operation after producing the selected signals.

To avoid problems that can be caused by asynchronous activities of the computing elements, the asynchronous activities are disabled. The functions of
30 the asynchronous activities are then performed when a selected signal is produced. For example, normal memory refresh functions are disabled and, in their place, burst memory refreshes are performed each time that a selected signal, such as an I/O request or a quantum interrupt, is
35 produced.

- 7 -

The invention also features a method of producing fault resilient or fault tolerant computers by designating a first processor as a computing element, designating a second processor as a controller, and

5 connecting the computing element and the controller to produce a modular pair. Thereafter, at least two modular pairs are connected to produce a fault resilient or fault tolerant computer. The processors used for the computing elements need not be identical to each other, but

10 preferably they all perform each instruction of their instruction sets in the same number of cycles as are taken by the other processors. Typically, industry standard processors are used in implementing the computing elements and the controllers. For disaster

15 tolerance, at least one of the modular pairs can be located remotely from the other modular pairs. The controllers and computing elements are each able to run unmodified industry standard operating systems and applications. In addition, the controllers are able to

20 run a first operating system while the computing elements simultaneously run a second operating system.

I/O fault resilience is obtained by connecting redundant I/O devices to at least two modular pairs and transmitting at least identical I/O write requests and

25 data to the redundant I/O devices. While I/O read requests need only be transmitted to one of the I/O devices, identical I/O read requests may be transmitted to more than one of the I/O devices to verify data integrity. When redundant I/O devices are connected to

30 three or more modular pairs, transmission of identical I/O requests allows identification of a faulty I/O device.

In another aspect, generally, the invention features isolating I/O requests from computing operations

35 in a computer through use of I/O redirection. Typically,

- 8 -

I/O devices are accessed either through low level I/O requests or by directly addressing the I/O devices. Low level I/O requests include requests to the system's basic input output system (e.g., BIOS), boot firmware requests, boot software requests, and requests to the system's physical device driver software. When a computing element issues a low level I/O request, the invention features using software to redirect the I/O requests to an I/O processor. When the computing element directly addresses the physical I/O devices, the invention features providing virtual I/O devices that simulate the interfaces of physical I/O devices. Directly addressed I/O requests are intercepted and provided to the virtual I/O devices. Periodically, the contents of the virtual I/O devices are transmitted to the I/O processor(s) as I/O requests. At the I/O processor(s), the transmitted contents of the virtual I/O devices are provided to the physical I/O devices. After the requested I/O operations are performed, the results of the operations, if any, are returned to the computing elements as responses to the I/O requests. Typically, the virtual I/O devices include a virtual keyboard and a virtual display.

The invention also features detecting and diagnosing faults in a computer system that includes at least two controllers that are connected to each other and to at least two computing elements, and at least two computing elements that are each connected to at least two of the controllers. Each computing element produces data and generates a value, such as an error checking code, that relates to the data. Each computing element then transmits the data, along with its corresponding value, to the at least two controllers to which it is connected. When the controllers receive the data and associated values, they transmit the values to the other controllers. Each controller then performs computations

- 9 -

on the values corresponding to each computing element and the values corresponding to each controller. If the results of the computations on the values corresponding to each controller are equal, and the results of the
5 computations on the values corresponding to each computing element are equal, then no fault exists. Otherwise, a fault exists. In some instances, the computation may be a simple bit by bit comparison.

When a fault exists, fault diagnosis is attempted
10 by comparing, for each one of the computing element, all of the values corresponding to the one computing element. If the values corresponding to each computing element match for each computing element, but mismatch for different computing elements, then one of the computing
15 elements is faulty. If the values corresponding to only one of the computing elements mismatch, then a path to that computing element is faulty. If the values corresponding to multiple computing elements mismatch, then the controller that is connected to the mismatching
20 computing elements is faulty. Once identified, the faulty element is disabled.

A system according to the invention can restore itself to full capability after a faulty element (that is, a CE, an IOP, a storage device, etc.) is repaired.
25 The system does so by transferring the state of an active element to the repaired element and, thereafter, reenabling the repaired element. Inactive or repaired processors are activated by transferring the operational state of an active processor or processors to the
30 inactive processor through a controller. When the inactive processor is a computing element, the operational state of an active computing element (or elements) is transferred through a controller. When the inactive processor is a controller, the operating state
35 of an active controller is directly transferred. The

- 10 -

transfer can occur either when system operation is paused or as a background process.

This recovery capability can also be used to provide on-line upgrades of hardware, software, or both by causing a processor of the system to fail by, for example, turning it off. The upgrade is then performed by either replacing or modifying the disabled processor. The upgraded processor is then turned on and reactivated as discussed above.

10 The invention also features a single controller, dual computing element system in which a controller is connected to two computing elements. In this computer system, I/O operations by the computing elements are intercepted and redirected to the controller. Typically,
15 the controller and the two computing elements each include an industry standard motherboard, and are each able to run unmodified industry standard operating systems and applications. In addition, the controller is able to run a first operating system while the computing
20 elements simultaneously run a second operating system.

The single controller system can be expanded to include a second controller connected both to the first controller and to the two computing elements. For purposes of providing limited disaster resilience, the
25 first controller and one of the computing elements can be placed in a location remote from the second controller and the other computing element, and can be connected to the second controller and the other computing element by a communications link.

30 For improved availability and performance, the dual controller, dual computing element system can be connected to an identical second system. The two systems then run a distributed computing environment in which one of the systems runs a first portion of a first

- 11 -

application and the other system runs either a second application or a second portion of the first application.

In another embodiment, the invention features a computer system that includes three controllers connected
5 to each other and three computing elements that are each connected to different pairs of the three controllers. This system, like the other systems, also features intercepting I/O operations by the computing elements and redirecting them to the controllers for processing. For
10 disaster resilience, the first controller and one of the computing elements are placed in a location remote from the remaining controllers and computing elements, or each controller/computing element pair is placed in a different location.

15 A disaster tolerant system is created by connecting at least two of the three controller systems described above. The three controller systems are placed in remote locations and connected by a communications link.

20 Brief Description of the Drawings

Fig. 1 is a block diagram of a partially fault resilient system.

Fig. 2 is a block diagram of system software of the system of Fig. 1.

25 Fig. 3 is a flowchart of a procedure used by an IOP Monitor of the system software of Fig. 2.

Fig. 4 is a block diagram of an IPI module of the system of Fig. 1.

Fig. 5 is a state transition table for the system
30 of Fig. 1.

Fig. 6 is a block diagram of a fault resilient system.

Fig. 7 is a block diagram of a distributed fault resilient system.

- 12 -

Fig. 8 is a block diagram of a fault tolerant system.

Fig. 9 is flowchart of a fault diagnosis procedure used by IOPs of the system of Fig. 8.

5 Fig. 10 is a block diagram of a disaster tolerant system.

Description of the Preferred Embodiments

As illustrated in Fig. 1, a fault resilient system 10 includes an I/O processor ("IOP") 12 and two computing elements ("CEs") 14a, 14b (collectively referred to as CEs 14). Because system 10 includes only a single IOP 12 and therefore cannot recover from a failure in IOP 12, system 10 is not entirely fault resilient.

IOP 12 includes two inter-processor interconnect 15 ("IPI") modules 16a, 16b that are connected, respectively, to corresponding IPI modules 18a, 18b of CEs 14 by cables 20a, 20b. IOP 12 also includes a processor 22, a memory system 24, two hard disk drives 26, 28, and a power supply 30. Similarly, each CE 14 20 includes a processor 32, a memory system 34, and a power supply 36. Separate power supplies 36 are used to ensure fault resilience in the event of a power supply failure. Processors 32a, 32b are "identical" to each other in that, for every instruction, the number of cycles 25 required for processor 32a to perform an instruction is identical to the number of cycles required for processor 32b to perform the same instruction. In the illustrated embodiment, system 10 has been implemented using standard Intel 486 based motherboards for processors 22, 32 and 30 four megabytes of memory for each of memory systems 24, 34.

IOP 12 and CEs 14 of system 10 run unmodified operating system and applications software, with hard drive 26 being used as the boot disk for the IOP and hard 35 drive 28 being used as the boot disk for CEs 14. In

- 13 -

truly fault resilient or fault tolerant systems that include at least two IOPs, each hard drive would also be duplicated.

In the illustrated embodiment, the operating system for IOP 12 and CEs 14 is DOS. However, other operating systems can also be used. Moreover, IOP 12 can run a different operating system from the one run by CEs 14. For example, IOP 12 could run Unix while CEs 14 run DOS. This approach is advantageous because it allows CEs 14 to access peripherals from operating systems that do not support the peripherals. For example, if CEs 14 were running an operating system that did not support CD-ROM drives, and IOP 12 were running one that did, CEs 14 could access the CD-ROM drive by issuing I/O requests identical to those used to, say, access a hard drive. IOP 12 would then handle the translation of the I/O request to one suitable for accessing the CD-ROM drive.

Referring also to Fig. 2, system 10 includes specialized system software 40 that controls the booting and synchronization of CEs 14, disables local time in CEs 14, redirects all I/O requests from CEs 14 to IOP 12 for execution, and returns the results of the I/O requests, if any, from IOP 12 to CEs 14.

System software 40 includes two sets of IPI BIOS 42 that are ROM-based and are each located in the IPI module 18 of a CE 14. IPI BIOS 42 are used in bootup and synchronization activities. When a CE 14 is booted, IPI BIOS 42 replaces the I/O interrupt addresses in the system BIOS interrupt table with addresses that are controlled by CE Drivers 44. The interrupt addresses that are replaced include those corresponding to video services, fixed disk services, serial communications services, keyboard services, and time of day services.

IPI BIOS 42 also disables normal memory refreshing to ensure that memory refreshing, which affects the

- 14 -

number of cycles during which a CE 14 is actually processing, is controlled by system software 40. Memory refreshing is required to maintain memory integrity. In known refreshing methods, memory is refreshed
5 periodically, with one block of memory being refreshed at the end of each refresh period. The duration of the refresh period is selected so that the entire memory is refreshed within the memory's refresh limit. Thus, for example, if a memory has 256 blocks and an 8 ms refresh
10 limit, then the refresh period is $31.25 \mu s$ ($8 \text{ ms} / 256$).

In the described embodiment, IPI BIOS 42 disables memory refreshing by placing a counter used in the Intel 486 motherboard to control memory refreshing in a mode that requires a gate input to the counter to change in
15 order to increment. Because the gate input is typically connected to the power supply, the gate input never changes and the counter is effectively disabled.

Two CE Drivers 44 of system software 40 handle memory refreshing by burst refreshing multiple blocks of
20 memory each time that an I/O request or quantum interrupt is generated. CE Drivers 44 are stored on CE boot disk 28 and are run by CEs 14. In addition to performing burst memory refreshes, CE Drivers 44 intercept I/O requests to the system BIOS and redirects them through
25 IPI modules 18 to IOP 12 for execution. CE Drivers 44 also respond to interrupt requests from IPI modules 18, disable the system clock, and, based on information supplied by IOP Monitor 48, control the time of day of CEs 14.

30 An IOP Driver 46 that is located on IOP boot disk 26 and is run by IOP 12 handles I/O requests from CEs 14 by redirecting them to an IOP Monitor 48 for processing and transmitting the results from IOP Monitor 48 to CEs 14. IOP Driver 46 communicates with CE drivers 44 using
35 a packet protocol.

- 15 -

IOP Monitor 48 is located on IOP boot disk 26 and is run by IOP 12. IOP Monitor 48 controls system 10 and performs the actual I/O requests to produce the results that are transmitted by IOP Driver 46 to CEs 14.

5 System software 40 also includes console software 49 that runs on IOP 12 and provides for user control of system 10. Using console software 49, a user can reset, boot, or synchronize a CE 14. The user can also set one or both of CEs 14 to automatically boot (autoboot) and/or
10 automatically synchronize (autosync) after being reset or upon startup. The ability to control each CE 14 is useful both during normal operation and for test purposes. Using console software 49, the user can also place system 10 into either an integrity mode in which
15 IOP Monitor 48 shuts down both CEs 14 when faced with a miscompare error, a first availability mode in which IOP Monitor 48 disables CE 14a when faced with a miscompare error, or a second availability mode in which IOP Monitor 48 disables CE 14b when faced with a miscompare error.
20 Finally, console software 49 allows the user to request the status of system 10. In an alternative embodiment, console software 49 could be implemented using a separate processor that communicates with IOP 12.

Each CE 14 runs a copy of the same application and
25 the same operating system as that run by the other CE 14. Moreover, the contents of memory systems 34a and 34b are the same, and the operating context of CEs 14 are the same at each synchronization time. Thus, IOP Monitor 48 should receive identical sequences of I/O requests from
30 CEs 14.

As shown in Fig. 3, IOP Monitor 48 processes and monitors I/O requests according to a procedure 100. Initially, IOP Monitor 48 waits for an I/O request from one of CEs 14 (step 102). Upon receiving an I/O request
35 packet from, for example, CE 14b, IOP Monitor 48 waits

- 16 -

for either an I/O request from CE 14a or for the expiration of a timeout period (step 104). Because system 10 uses the DOS operating system, which halts execution of an application while an I/O request is being
5 processed, IOP Monitor 48 is guaranteed not to receive an I/O request from CE 14b while waiting (step 104) for the I/O request from the CE 14a.

Next, IOP Monitor 48 checks to determine whether the timeout period has expired (step 106). If not (that
10 is, an I/O request packet from CE 14a has arrived), IOP Monitor 48 compares the checksums of the packets (step 108), and, if the checksums are equal, processes the I/O request (step 110). After processing the I/O request, IOP Monitor 48 issues a request to the system BIOS of IOP
15 12 for the current time of day (step 112).

After receiving the time of day, IOP Monitor 48 assembles an IPI packet that includes the time of day and the results, if any, of the I/O request (step 114) and sends the IPI packet to IOP Driver 46 (step 116) for
20 transmission to CEs 14. When CEs 14 receive the IPI packet, they use the transmitted time of day to update their local clocks which, as already noted, are otherwise disabled.

As required by DOS, execution in CEs 14 is
25 suspended until IOP Monitor 48 returns the results of the I/O request through IOP Driver 46. Because, before execution is resumed, the times of day of both CEs 14 are updated to a common value (the transmitted time of day from the IPI packet), the CEs 14 are kept in time
30 synchronization with the transmitted time of day being designated the meta time. If a multitasking operating system were employed, execution in CEs 14 would not be suspended while IOP Monitor 48 performed the I/O request. Instead, processing in CEs 14 would be suspended only
35 until receipt of an acknowledgement indicating that IOP

- 17 -

Monitor 48 has begun processing the I/O request (step 110). The acknowledgement would include the time of day and would be used by CEs 14 to update the local clocks.

After sending the IPI packet to IOP Driver 46, IOP
5 Monitor 48 verifies that both of CEs 14 are online (step 118), and, if so, waits for another I/O request from one of CEs 14 (step 102).

If the timeout period has expired (step 106), IOP
Monitor 48 disables the CE 14 that failed to respond
10 (step 119) and processes the I/O request (step 110).

If there is a miscompare between the checksums of the packets from CEs 14 (step 108), IOP Monitor 48 checks to see if system 10 is operating in an availability mode or an integrity mode (step 120).

15 If system 10 is operating in an availability mode, IOP Monitor 48 disables the appropriate CE 14 based on the selected availability mode (step 122), and processes the I/O request (step 110). Thereafter, when IOP Monitor 48 checks whether both CEs 14 are online (step 118), and
20 assuming that the disabled CE 14 has not been repaired and reactivated, IOP Monitor 48 then waits for an I/O request from the online CE 14 (step 124). Because system 10 is no longer fault resilient, when an I/O request is received, IOP Monitor 48 immediately processes the I/O
25 request (step 110).

If system 10 is operating in an integrity mode when a miscompare is detected, IOP Monitor 48 disables both CEs 14 (step 126) and stops processing (step 128).

Referring again to Figs. 1 and 2, when the
30 application or the operating system of, for example, CE 14a makes a non-I/O call to the system BIOS, the system BIOS executes the request and returns the results to the application without invoking system software 40.

However, if the application or the operating system makes
35 an I/O BIOS call, CE Driver 44a intercepts the I/O

- 18 -

request. After intercepting the I/O request, CE Driver 44a packages the I/O request into an IPI packet and transmits the IPI packet to IOP 12.

When IPI module 16a of IOP 12 detects transmission of an IPI packet from CE 14a, IPI module 16a generates an interrupt to IOP Driver 16. IOP Driver 46 then reads the IPI packet.

As discussed above, IOP Monitor 48 responds to the IPI packet from CE 14a according to procedure 100. As also discussed, assuming that there are no hardware faults, IOP Driver 46 eventually transmits an IPI packet that contains the results of the I/O request and the time of day to CEs 14.

IPI modules 18 of CEs 14 receive the IPI packet from IOP 12. CE Drivers 44 unpack the IPI packet, update the time of day of CEs 14, and return control of CEs 14 to the application or the operating system running on CEs 14.

If no I/O requests are issued within a given time interval, the IPI module 18 of a CE 14 generates a so-called quantum interrupt that invokes the CE Driver 44 of the CE 14. In response, the CE Driver 44 creates a quantum interrupt IPI packet and transmits it to IOP 12. IOP Monitor 48 treats the quantum interrupt IPI packet as an IPI packet without an I/O request. Thus, IOP Monitor 48 detects the incoming quantum interrupt IPI packet (step 102 of Fig. 3) and, if a matching quantum interrupt IPI packet is received from the other CE 14 (steps 104, 106, and 108 of Fig. 3), issues a request to the system BIOS of IOP 12 for the current time of day (step 112 of Fig. 3). IOP Monitor 48 then packages the current time of day into a quantum response IPI packet (step 114 of Fig. 3) that IOP Driver 46 then sends to CEs 14 (step 116 of Fig. 3). CE Drivers 44 respond to the quantum response IPI packet by updating the time of day and

- 19 -

returning control of CEs 14 to the application or the operating system running on CEs 14.

If IOP Monitor 48 does not receive a quantum interrupt IPI package from the other CE 14 within a predefined timeout period (step 106 of Fig. 3), IOP Monitor 48 responds by disabling the non-responding CE 14.

As shown in Fig. 1, IPI modules 16, 18 and cables 20 provide all of the hardware necessary to produce a fault resilient system from the standard Intel 486 based motherboards used to implement processors 22, 32. An IPI module 16 and an IPI module 18, which are implemented using identical boards, each perform similar functions.

As illustrated in Fig. 4, an IPI module 18 includes a control logic 50 that communicates I/O requests and responses between the system bus of a processor 32 of a CE 14 and a parallel interface 52 of IPI module 18. Parallel interface 52, in turn, communicates with the parallel interface of an IPI module 16 through a cable 20. Parallel interface 52 includes a sixteen bit data output port 54, a sixteen bit data input port 56, and a control port 58. Cable 20 is configured so that data output port 54 is connected to the data input port of the IPI module 16, data input port 56 is connected to the data output port of the IPI module 16, and control port 58 is connected to the control port of the IPI module 16. Control port 58 implements a handshaking protocol between IPI module 18 and the IPI module 16.

Control logic 50 is also connected to an IPI BIOS ROM 60. At startup, control logic 50 transfers IPI BIOS 42 (Fig. 2), the contents of IPI BIOS ROM 60, to processor 32 through the system bus of processor 32.

A QI counter 62, also located on IPI module 18, generates quantum interrupts as discussed above. QI

- 20 -

counter 62 includes a clock input 64 that is connected to the system clock of processor 32 and a gate input 66 that is connected to control logic 50. Gate input 66 is used to activate and reset the counter value of QI counter 62.

- 5 When activated, QI counter 62 decrements the counter value by one during each cycle of the system clock of processor 32. When the counter value reaches zero, QI counter 62 generates a quantum interrupt that, as discussed above, activates CE Driver 44 (Fig. 2).
- 10 CE Driver 44 deactivates QI counter 62 at the beginning of each I/O transaction. CE Driver 44 deactivates QI counter 62 by requesting an I/O write at a first address, known as the QI deactivation address. Control logic 50 detects the I/O write request and
- 15 deactivates QI counter 62 through gate input 66. Because this particular I/O write is for control purposes only, control logic 50 does not pass the I/O write to parallel interface 52. At the conclusion of each I/O transaction, CE Driver 44 resets and activates QI counter 62 by
- 20 requesting an I/O write to a second address, known as the QI activation address. Control logic 50 responds by resetting and activating QI counter 62.

- In an alternative approach, quantum interrupts are generated through use of debugging or other features
- 25 available in processor 32. Some commonly available processors include debugging or trap instructions that trap errors by transferring control of the processor to a designated program after the completion of a selected number of instructions following the trap instruction.
- 30 In this approach, each time that CE Driver 44 returns control of processor 32 to the application or operating system, CE Driver 44 issues a trap instruction to indicate that control of processor 32 should be given to CE Driver 44 upon completion of, for example, 300
- 35 instructions. After processor 32 completes the indicated

- 21 -

300 instructions, the trap instruction causes control of processor 32 to be returned to CE Driver 44. In the event that an I/O request activates CE Driver 44 prior to completion of the indicated number of instructions, CE
5 Driver 44 issues an instruction that cancels the trap instruction.

IPI Module 18 is also used in activating an offline CE 14. As discussed below, before an offline CE 14 is activated, the contents of the memory system 34 of
10 the active CE 14 are copied into the memory system 34 of the offline CE 14. To minimize the effects of this copying on the active CE 14, the processor 32 of the active CE 14 is permitted to continue processing and the memory is copied only during cycles in which the system
15 bus of the processor 32 of the active CE 14 is not in use.

To enable processor 32 to continue processing while the memory is being copied, IPI module 18 accounts for memory writes by the processor 32 to addresses that
20 have already been copied to the offline CE 14. To do so, control logic 50 monitors the system bus and, when the processor 32 writes to a memory address that has already been copied, stores the address in a FIFO 68. When the memory transfer is complete, or when FIFO 68 is full, the
25 contents of memory locations associated with the memory addresses stored in FIFO 68 are copied to the offline CE 14 and FIFO 68 is emptied. In other approaches, FIFO 68 is modified to store both memory addresses and the contents of memory locations associated with the
30 addresses, or to store the block addresses of memory blocks to which memory addresses being written belong.

IPI module 18 also handles non-BIOS I/O requests. In some computer systems, the BIOS is too slow to effectively perform I/O operations such as video display.
35 As a result, some less structured or less disciplined

- 22 -

operating systems, such as DOS or UNIX, allow applications to circumvent the BIOS and make non-BIOS I/O requests by directly reading from or writing to the addresses associated with I/O devices. These non-BIOS
5 I/O requests, which cannot be intercepted by changing the system interrupt table, as is done in connection with, for example, I/O disk reads and writes, are problematic for a system in which synchronization requires tight control of the I/O interface.

10 To remedy this problem, and to assure that even non-BIOS I/O requests can be isolated and managed by IOP 12, IPI module 18 includes virtual I/O devices that mimic the hardware interfaces of physical I/O devices. These virtual I/O devices include a virtual display 70 and a
15 virtual keyboard 72. As needed, other virtual I/O devices such as a virtual mouse or virtual serial and parallel ports could also be used.

In practice, control logic 50 monitors the system bus for read or write operations directed to addresses
20 associated with non-BIOS I/O requests to system I/O devices. When control logic 50 detects such an operation, control logic 50 stores the information necessary to reconstruct the operation in the appropriate virtual device. Thus, for example, when control logic 50
25 detects a write operation directed to an address associated with the display, control logic 50 stores the information necessary to reconstruct the operation in virtual display 70. Each time that a BIOS I/O request or a quantum interrupt occurs, CE Driver 44 scans the
30 virtual I/O devices and, if the virtual devices are not empty, assembles the information stored in the virtual devices into an IPI packet and transmits the IPI packet to IOP 12. IOP 12 treats the packet like a BIOS I/O request using procedure 100 discussed above. When
35 control logic 50 detects a read addressed to a virtual

- 23 -

I/O device, control logic 50 assembles the read request into an IPI packet for handling by IOP 12. IOP 12 treats the IPI packet like a standard BIOS I/O request.

Referring to Fig. 5, each CE 14 always operates in one of eight states and, because there are only a limited number of permissible state combinations, system 10 always operates in one of fourteen states. The major CE operating states are OFFLINE, RTB (ready to boot), BOOTING, ACTIVE, RTS (ready to sync), WAITING, M_SYNC, (synchronizing as master), and S_SYNC (synchronizing as slave). IOP Monitor 48 changes the operating states of CEs 14 based on the state of system 10 and user commands from console software 49. Through console software 49, a user can reset a CE 14 at any time. Whenever the user resets a CE 14, or a fault occurs in the CE 14, IOP Monitor 48 changes the state of the CE 14 to OFFLINE.

At startup, system 10 is operating with both CEs 14 OFFLINE (state 150). System 10 operates in the upper states of Fig. 5 (states 152-162) when CE 14a becomes operational before CE 14b and in the lower states (states 166-176) when CE 14b is the first to become operational. If CEs 14 become operational simultaneously, the first operational CE 14 to be recognized by IOP Monitor 48 is treated as the first to become operational.

When a CE 14 indicates that it is ready to boot by issuing a boot request, the state of the CE 14 changes to RTB if the CE 14 is not set to autoboot or to BOOTING if the CE 14 is set to autoboot. For example, if CE 14a issues a boot request when both CEs 14 are OFFLINE, and CE 14a is not set to autoboot, then the state of CE 14a changes to RTB (state 152). Thereafter, IOP Monitor 48 waits for the user, through console software 49, to boot CE 14a. When the user boots CE 14a, the state of CE 14a changes to BOOTING (state 154). If the user resets CE 14a, the state of CE 14a changes to OFFLINE (state 150).

- 24 -

If both CEs 14 are OFFLINE when CE 14a issues a boot request, and CE 14a is set to autoboot, the state of CE 14a changes to BOOTING (state 154). If CE 14a boots successfully, the state of CE 14a changes to ACTIVE
5 (state 156).

When CE 14a is ACTIVE, and CE 14b issues a boot request, or if CE 14b had issued a boot request while the state of CE 14a was transitioning from OFFLINE to ACTIVE (states 152-156), the state of CE 14b changes to RTS
10 (state 158) if CE 14b is set to autosync and otherwise to WAITING (state 160). If the state of CE 14b changes to RTS (state 158), IOP Monitor waits for the user to issue a synchronize command to CE 14b. When the user issues such a command, the state of CE 14b changes to WAITING
15 (state 160).

Once CE 14b is WAITING, IOP Monitor 48 copies the contents of memory system 34a of CE 14a into memory system 34b of CE 14b. Once the memory transfer is complete, IOP Monitor 48 waits for CE 14a to transmit a
20 quantum interrupt or I/O request IPI packet. Upon receipt of such a packet, IOP Monitor 48 changes the state of CE 14a to M_SYNC and the state of CE 14b to S_SYNC (state 162), and synchronizes the CEs 14. This synchronization includes responding to any memory changes
25 that occurred while IOP Monitor 48 was waiting for CE 14a to transmit a quantum interrupt or I/O request IPI packet. Upon completion of the synchronization, the states of the CEs 14 both change to ACTIVE (state 164) and system 10 is deemed to be fully operational.

30 In an alternative implementation, IOP Monitor 48 does not wait for memory transfer to complete before changing the state of CE 14a to M_SYNC and the state of CE 14b to S_SYNC (state 162). Instead, IOP Monitor 48 makes this state change upon receipt of an IPI packet

- 25 -

from CE 14a and performs the memory transfer as part of the synchronization process.

Similar state transitions occur when CE 14b is the first CE 14 to issue a boot request. Thus, assuming that
5 CE 14b is not set to autoboot, CE 14b transitions from OFFLINE (state 150) to RTC (state 166) to BOOTING (state 168) to ACTIVE (state 170). Similarly, once CE 14b is ACTIVE, and assuming that CE 14a is not set to autosync, CE 14a transitions from OFFLINE (state 170) to RTS (state
10 172) to WAITING (state 174) to S_SYNC (state 176) to ACTIVE (state 164).

In other embodiments of the invention, for example, referring to Fig. 6, a fault resilient system 200 includes two IOPs 202 and two CEs 204. Each CE 204
15 is connected, through an IPI card 206 and a cable 208, to an IPI card 210 of each IOP 202. IOPs 202 are redundantly connected to each other through IPI cards 210 and cables 212. Because every component of system 200 has a redundant backup component, system 200 is entirely
20 fault resilient. In an alternative approach, cables 208 and 210 could be replaced by a pair of local area networks to which each IOP 202 and CE 204 would be connected. Indeed, local area networks can always be substituted for cable connections.

25 System 200 is operating system and application software independent in that it does not require modifications of the operating system or the application software to operate. Any single piece of hardware can be upgraded or repaired in system 200 with no service
30 interruption. Therefore, by sequentially replacing each piece of hardware and allowing system 200 to resynchronize after each replacement, the hardware of system 200 can be replaced in its entirety without service interruption. Similarly, software on system 200
35 can be upgraded with minimal service interruption (that

- 26 -

is, during the software upgrade, the application will become unavailable for an acceptable period of time such as two seconds). Also, disaster tolerance for purposes of availability can be obtained by placing each IOP/CE pair in a separate location and connecting the pairs through a communications link.

Referring to Fig. 7, a distributed, high performance, fault resilient system 220 includes two systems 200, the IOPs 202 of which are connected to each other, through IPI modules, by cables 222. System 220 uses distributed computing environment software to achieve high performance by running separate portions of an application on each system 200. System 220 is fault tolerant and offers the ability to perform both hardware and software upgrades without service interruption.

Referring to Fig. 8, a fault tolerant system 230 includes three IOPs (232, 234, and 236) and three CEs (238, 240, and 242). Through IPI modules 244 and cables 246, each IOP is connected to an IPI module 244 of each of the other IOPs. Through IPI modules 248 and cables 250, each CE is connected to an IPI module 244 of two of the IOPs, with CE 238 being connected to IOPs 232 and 234, CE 240 being connected to IOPs 232 and 236, and CE 242 being connected to IOPs 234 and 236. Like system 200, system 230 allows for hardware upgrades without service interruption and software upgrades with only minimal service interruption.

As can be seen from a comparison of Figs. 7 and 8, the CEs and IOPs of systems 200 and 230 are identically configured. As a result, upgrading a fault resilient system 200 to a fault tolerant system 230 does not require any replacement of existing hardware and entails the simple procedure of adding an additional CE/IOP pair, connecting the cables, and making appropriate changes to the system software. This modularity is an important

- 27 -

feature of the paired modular redundant architecture of the invention.

Because the components of system 230 are triply redundant, system 230 is more capable of identifying the source of a hardware fault than is system 10. Thus, while system 10 simply disables one or both of CEs 14 when an error is detected, system 230 offers a higher degree of fault diagnosis.

Referring to Fig. 9, each IOP (232, 234, 236) of system 230 performs fault diagnosis according to a procedure 300. Initially, each IOP (232, 234, 236) checks for major faults such as power loss, broken cables, and nonfunctional CEs or IOPs using well known techniques such as power sensing, cable sensing, and protocol timeouts (step 302). When such a fault is detected, each IOP disables the faulty device or, if necessary, the entire system.

After checking for major faults, each IOP waits to receive IPI packets (that is, quantum interrupts or I/O requests) from the two CEs to which the IOP is connected (step 304). Thus, for example, IOP 232 waits to receive IPI packets from CEs 238 and 240. After receiving IPI packets from both connected CEs, each IOP transmits the checksums ("CRCs") of those IPI packets to the other two IOPs and waits for receipt of CRCs from the other two IOPs (step 306).

After receiving the CRCs from the other two IOPs, each IOP generates a three by three matrix in which each column corresponds to a CE, each row corresponds to an IOP, and each entry is the CRC received from the column's

- 28 -

CE by the row's IOP (step 308). Thus, for example, IOP 232 generates the following matrix:

		CE 238	CE 240	CE 242
5	IOP 232	CRC	CRC	X
	IOP 234	CRC	X	CRC
	IOP 236	X	CRC	CRC

After generating the matrix, IOP 232 sums the entries in each row and each column of the matrix. If the three row sums are equal and the three column sums are equal (step 10 310), then there is no fault and IOP 232 checks again for major faults (step 302).

If either the three rows' sums or the three columns' sums are unequal (step 310), then IOP 232 compares the CRC entries in each of the columns of the 15 matrix. If the two CRC entries in each column match (step 312), then IOP 232 diagnoses that a CE failure has occurred and disables the CE corresponding to the column for which the sum does not equal the sums of the other columns (step 314).

20 If the CRC entries in one or more of the matrix columns do not match (step 312), then IOP 232 determines how many of the columns include mismatched entries. If the matrix includes only one column with mismatched entries (step 315), then IOP 232 diagnoses that the path 25 between the IOP corresponding to the matrix row sum that is unequal to the other matrix row sums and the CE corresponding to the column having mismatched entries has failed and disables that path (step 316). For purposes of the diagnosis, the path includes the IPI module 244 in 30 the IOP, the IPI module 248 in the CE, and the cable 250.

If the matrix includes more than one column with mismatched entries (step 314), then IOP 232 confirms that one matrix row sum is unequal to the other matrix row sums, diagnoses an IOP failure, and disables the IOP 35 corresponding to the matrix row sum that is unequal to the other matrix row sums (step 318).

- 29 -

If, after diagnosing and accounting for a CE failure (step 314), path failure (step 316), or IOP failure (step 318), IOP 232 determines that system 300 still includes sufficient non-faulty hardware to remain operational, IOP 232 checks again for major faults (step 302). Because system 230 is triply redundant, system 230 can continue to operate even after several components have failed. For example, to remain operating in an availability mode, system 230 only needs to have a single functional CE, a single functional IOP, and a functional path between the two.

Using procedure 300, each IOP (232, 234, 236) can correctly diagnose any single failure in a fully operational system 230 or in a system 230 in which one element (that is, a CE, an IOP, or a path) has previously been disabled. In a system 230 in which an element has been disabled, each IOP accounts for CRCs that are not received because of the disabled element by using values that appear to be correct in comparison to actually received CRCs.

Procedure 300 is not dependent on the particular arrangement of interconnections between the CEs and IOPs. To operate properly, procedure 300 only requires that the output of each CE be directly monitored by at least two IOPs. Thus, procedure 300 could be implemented in a system using any interconnect mechanism and does not require point to point connections between the CEs and IOPs. For example, the CEs and IOPs could be connected to at least two local area networks. In an alternative approach, instead of summing the CRC values in the rows and columns of the matrix, these values can be compared and those rows or columns in which the entries do not match can be marked with a match/mismatch indicator.

A simplified version of procedure 300 can be implemented for use in a system 200. In this procedure,

- 30 -

each IOP 202 of system 200 generates a two by two matrix in which each column corresponds to a CE 204 and each row corresponds to a IOP 202:

		CE 204	CE 204
5	IOP 202	CRC	CRC
	IOP 202	CRC	CRC

After generating the matrix, each IOP 202 attaches a mismatch indicator to each row or column in which the two entries are mismatched.

10 If there are no mismatch indicators, then system 200 is operating correctly.

If neither row and both columns have mismatch indicators, then an IOP 202 has faulted. Depending on the operating mode of system 200, an IOP 202 either
 15 disables another IOP 202 or shuts down system 200. The IOP 202 to be disabled is selected based on user supplied parameters similar to the two availability modes used in system 10.

If both rows and neither column have mismatch
 20 indicators, then a CE 204 has faulted. In this case, IOPs 202 respond by disabling a CE 204 if system 200 is operating in an availability mode or, if system 200 is operating in an integrity mode, shutting down system 200. If both rows and one column have mismatch indicators,
 25 then one of the paths between the IOPs 202 and the CE 204 corresponding to the mismatched column has failed. Depending on the operating mode of system 200, IOPs 202 either disable the CE 204 having the failed path or shut down system 200. If both rows and both column have
 30 mismatch indicators, then multiple faults exist and IOPs 202 shut down system 200.

If one row and both columns have mismatch indicators, then the IOP 202 corresponding to the mismatched row has faulted. Depending on the operating
 35 mode of system 200, the other IOP 202 either disables the faulty IOP 202 or shuts down system 200. If one row and

- 31 -

one column have mismatch indicators, then the path between the IOP 202 corresponding to the mismatched row and the CE 204 corresponding to the mismatched column has failed. Depending on the operating mode of system 200, 5 IOPs 202 either account for the failed path in future processing or shut down system 200.

Referring to Fig. 10, one embodiment of a disaster tolerant system 260 includes two fault tolerant systems 230 located in remote locations and connected by 10 communications link 262, such as Ethernet or fiber, and operating in meta time lockstep with each other. To obtain meta time lockstep, all IPI packets are transmitted between fault tolerant systems 230. Like system 220, system 260 allows for hardware and software 15 upgrades without service interruption.

As shown, the paired modular redundant architecture of the invention allows for varying levels of fault resilience and fault tolerance through use of CEs that operate asynchronously in real time and are 20 controlled by IOPs to operate synchronously in meta time. This architecture is simple and cost-effective, and can be expanded or upgraded with minimal difficulty.

What is claimed is:

- 32 -

1. A method of synchronizing at least two computing elements in a computer system including the at least two computing elements and at least one controller, wherein each of the computing elements have clocks that
5 operate asynchronously of the clocks of the other computing elements, said method comprising the steps of:
 selecting one or more signals from a set of signals produced by the computing elements;
 monitoring the computing elements to detect the
10 production of a selected signal by one of the computing elements;
 waiting for the production of a selected signal by each other computing element after detection of a selected signal by one of the computing elements;
15 transmitting equal time updates from the at least one controller to each of the computing elements after receipt of selected signals from all of the computing elements; and
 updating the clocks of the computing elements
20 based on the time updates.
2. The method of claim 1, further comprising the step of forming a fault resilient computer from the at least two computing elements and the at least one controller.
- 25 3. The method of claim 1, wherein said selecting step comprises the step of selecting I/O requests as the selected signals.
4. The method of claim 3, further comprising the steps of:
30 processing the I/O requests at the at least one controller to produce I/O responses; and

- 33 -

transmitting the time updates with the I/O responses from the at least one controller to the at least two computing elements.

5 5. The method of claim 1, wherein said selecting step comprises the step of selecting quantum interrupts and I/O requests as the selected signals.

6. The method of claim 1, wherein said selecting step comprises the step of selecting quantum interrupts as the selected signals.

10 7. The method of claim 6, further comprising the step of generating quantum interrupts in each computing element by counting clock cycles in the computing elements.

15 8. The method of claim 7, wherein the step of counting clock cycles includes counting the cycles of a selected one of a system clock, an I/O clock, and a bus clock.

9. The method of claim 7, further comprising the steps of:

20 loading a counter in each of the computing elements with a predetermined value;
 enabling the counter in each of the computing elements with an I/O request;
 decrementing the value of the counter during a
25 clock cycle in each of the computing elements; and
 signalling a quantum interrupt from a computing element when the value of the counter of the computing element reaches zero.

- 34 -

10. The method of claim 6, further comprising the step of generating quantum interrupts by counting executed instructions in each computing element.

11. The method of claim 6, further comprising the step of using debugging features of each computing element to generate quantum interrupts.

12. The method of claim 1, further comprising the step of maintaining, for each computing element, a list of the selected signals produced by the computing element, wherein the equal time updates are transmitted when the lists for each computing element include a common entry.

13. The method of claim 1, further comprising the steps of:

15 comparing the selected signals generated by the computing elements and data, if any, accompanying the selected signals, and

signalling that a fault has occurred if the selected signals or the accompanying data do not match.

14. The method of claim 1, further comprising the steps of:

stopping operation of each computing element after that computing element produces a selected signal, and

resuming operation of a computing element upon receipt by the computing element of a time update.

15. The method of claim 1, further comprising the step of continuing operation of a computing element after producing the selected signal.

- 35 -

16. The method of claim 1, further comprising the steps of:

disabling asynchronous activities of the computing elements; and

- 5 performing the functions of the asynchronous activities at a computing element when the computing element produces a selected signal.

17. The method of claim 16, wherein said disabling step comprises the step of disabling normal
10 memory refresh functions, and said performing step comprises the step of performing burst memory refreshes when said selected signal is produced.

18. The method of claim 17, wherein said disabling step further comprises the steps of:

- 15 placing a counter used in the normal memory refresh functions in a mode that requires an input value to a gate to change, and
connecting the gate to a fixed voltage.

19. A method of producing a fault resilient or
20 fault tolerant computer, comprising the steps of:

designating a first processor as a computing element;

designating a second processor as a controller;

connecting the computing element and the
25 controller to produce a modular pair;

connecting at least two modular pairs to produce a fault resilient or fault tolerant computer,

wherein each computing element performs all instructions in the same number of cycles as the other
30 computing elements.

- 36 -

20. The method of claim 19, wherein the first and second processors are industry standard processors.

21. The method of claim 19, further including the step of running industry standard operating systems and
5 applications on the at least two controllers and the at least two computing elements.

22. The method of claim 19, further including the steps of:

running a first operating system on the at least
10 two controllers; and

running a second operating system on the at least two computing elements.

23. The method of claim 19, further comprising the step of locating a modular pair remotely from the one
15 or more other modular pairs to provide disaster tolerance.

24. The method of claim 19, further comprising the steps of:

connecting a first I/O device to a first modular
20 pair;

connecting a second I/O device to a second modular pair, said second I/O device being redundant of the first I/O device; and

transmitting at least identical I/O write requests
25 and data to the first and second I/O devices.

25. The method of claim 24, further comprising the steps of:

connecting a third I/O device to a third modular pair, said third I/O device being redundant of the first
30 and second I/O devices; and

- 37 -

transmitting at least identical I/O write requests and data to the first, second, and third I/O devices.

26. The method of claim 19, further comprising the step of activating an inactive processor by
5 transferring the operational state of an active processor to the inactive processor through a controller.

27. The method of claim 26, further comprising the step of pausing processing by said computing elements during said transferring step.

10 28. The method of claim 26, further comprising the step of performing said transferring step as a background process without pausing processing by said computing elements.

29. The method of claim 19, further comprising
15 the step of upgrading a processor while said computing elements are processing by:
 disabling a processor to be upgraded;
 upgrading the disabled processor; and
 reactivating the upgraded processor by
20 transferring the operational state of an active processor to the upgraded processor through a controller.

30. The method of claim 19, further comprising the step of repairing a processor while said computing elements are processing by:
25 disabling a processor to be repaired;
 repairing the disabled processor; and
 reactivating the repaired processor by
transferring the operational state of an active processor to the repaired processor through a controller.

- 38 -

31. A method of isolating I/O requests from computing operations in a computer, comprising the steps of:

- 5 providing a virtual I/O unit that simulates the interface of a physical I/O device;
- intercepting an I/O request by a computing element that is addressed to the physical I/O device;
- providing the intercepted I/O request to the virtual I/O unit;
- 10 transmitting the contents of the virtual I/O unit to an I/O processor; and
- at the I/O processor, providing the transmitted contents of the virtual I/O device to the physical I/O device.

15 32. The method of claim 31, wherein said providing step includes providing a virtual keyboard.

33. The method of claim 31, wherein said providing step includes providing a virtual display.

20 34. The method of claim 31, further comprising the step of using the virtual I/O device to expose software errors caused by a software asynchrony that results in I/O race conditions.

35. The method of claim 31, further comprising the steps of:

- 25 intercepting a low level I/O request by a computing element;
- redirecting the intercepted low level I/O request to the I/O processor;
- at the I/O processor, performing the requested I/O
- 30 operation to produce I/O results; and

- 39 -

returning the I/O results to the computing element.

36. A method of detecting and diagnosing faults in a computer system that includes at least two computing
5 elements and at least two controllers, wherein each of the computing elements is connected to at least two of the controllers, and each controller is connected to at least two computing elements and to the other controllers, said method comprising the steps of:
10 producing data at each of the computing elements;
 generating a value at each of the computing elements that relates to the produced data;
 transmitting the data, along with the corresponding values, from each computing element to the
15 at least two connected controllers;
 transmitting the values received by each controller to the other controllers; and
 performing computations on the values corresponding to each computing element and the values
20 corresponding to each controller;
 wherein, when the results of the computations performed on the values corresponding to each controller are equal, and the results of the computations performed on the values corresponding to each computing element are
25 equal, no faults exist.

37. The method of claim 36, further comprising, when the results of the computations performed on the values corresponding to each computing element and the results of the computations performed on the values
30 corresponding to each controller are not equal, the steps of:

- 40 -

comparing, for each one of the computing elements,
all of the values corresponding to the one computing
element, and

designating one of the computing elements as
5 faulty when the values corresponding to each computing
element match for each computing element, but mismatch
for different computing elements.

38. The method of claim 36, further comprising,
when the results of the computations performed on the
10 values corresponding to each computing element and the
results of the computations performed on the values
corresponding to each controller are not equal, the steps
of:

comparing, for each one of the computing elements,
15 all of the values corresponding to the one computing
element, and

designating a connection to one of the computing
elements as faulty when the values corresponding only to
the one computing element mismatch.

20 39. The method of claim 36, further comprising,
when the results of the computations performed on the
values corresponding to each computing element and the
results of the computations performed on the values
corresponding to each controller are not equal, the steps
25 of:

comparing, for each one of the computing elements,
all of the values corresponding to the one computing
element, and

when the values corresponding to two or more of
30 the computing elements mismatch, designating the
controller connected to the two or more computing
elements as faulty.

- 41 -

40. A computer system, comprising:
a controller,
a first computing element connected to the controller,
5 a second computing element connected to the controller,
means for intercepting I/O operations by the first and second computing elements, and
means for transmitting the intercepted I/O
10 operations to the controller,
wherein the first computing element performs each instruction of its instruction set in the same number of cycles as the second computing element takes to perform said instruction.
- 15 41. The computer system of claim 40, wherein the controller and the first and second computing elements each include an industry standard motherboard.
- 20 42. The computer system of claim 40, further comprising a second controller connected to the first controller and to the first and second computing elements.
- 25 43. The computer system of claim 42, wherein the first controller and the first computing element are located in a first location and the second controller and the second computing element are located in a second location, and further comprising a communications link connecting the first controller to the second controller, the first controller to the second computing element, and the second controller to the first computing element.
- 30 44. The computer system of claim 42, further comprising:

- 42 -

a third controller;
a fourth controller connected to the third
controller;
a third computing element connected to the third
5 controller and the fourth controller;
a fourth computing element connected to the third
controller and the fourth controller;
means for connecting the third and fourth
controllers to the first and second controller; and
10 means for distributing computing tasks between the
computing elements, wherein the first and second
computing elements perform a first set of computing tasks
and the third and fourth computing elements perform a
second set of computing tasks
15 wherein the third and fourth computing elements
perform each instruction of their instruction sets in the
same number of cycles as the first and second computing
elements take to perform said instruction.

45. The computer system of claim 42, wherein the
20 first controller and the first computing element are
remotely located from the second controller and the
second computing element to provide disaster tolerance.

46. The computer system of claim 40, wherein each
of said first and second computing elements further
25 comprises means for generating a quantum interrupt.

47. A computer system, comprising:
a first controller;
a second controller connected to the first
controller;
30 a third controller connected to the first and
second controllers;

- 43 -

a first computing element connected to the first and second controllers;

a second computing element connected to the second and third controllers; and

5 a third computing element connected to the first and third controllers.

48. The computer system of claim 47, wherein the first controller and the first computing element are remotely located from the other controllers and computing
10 elements to provide disaster tolerance.

49. The computer system of claim 47, further comprising:

means for intercepting I/O operations by the first computing element;

15 means for transmitting the intercepted I/O from the first computing element to the first and second controllers;

means for intercepting I/O operations by the second computing element;

20 means for transmitting the intercepted I/O from the second computing element to the second and third controllers;

means for intercepting I/O operations by the third computing element; and

25 means for transmitting the intercepted I/O from the third computing element to the first and third controllers.

50. The computer system of claim 47, further comprising:

30 a fourth controller;

a fifth controller connected to the fourth controller;

- 44 -


a sixth controller connected to the fourth and fifth controllers;

a fourth computing element connected to the fourth and fifth controllers;

5 a fifth computing element connected to the fifth and sixth controllers;

a sixth computing element connected to the fourth and sixth controllers; and

a communications link for connecting the first,
10 second, and third controllers to the fourth, fifth, and sixth controllers, wherein the first, second, and third controllers, and the first, second, and third computing elements are in a first location and the fourth, fifth, and sixth controllers, and the fourth, fifth, and sixth
15 computing elements are in a second location.






FIG. 2

3/9

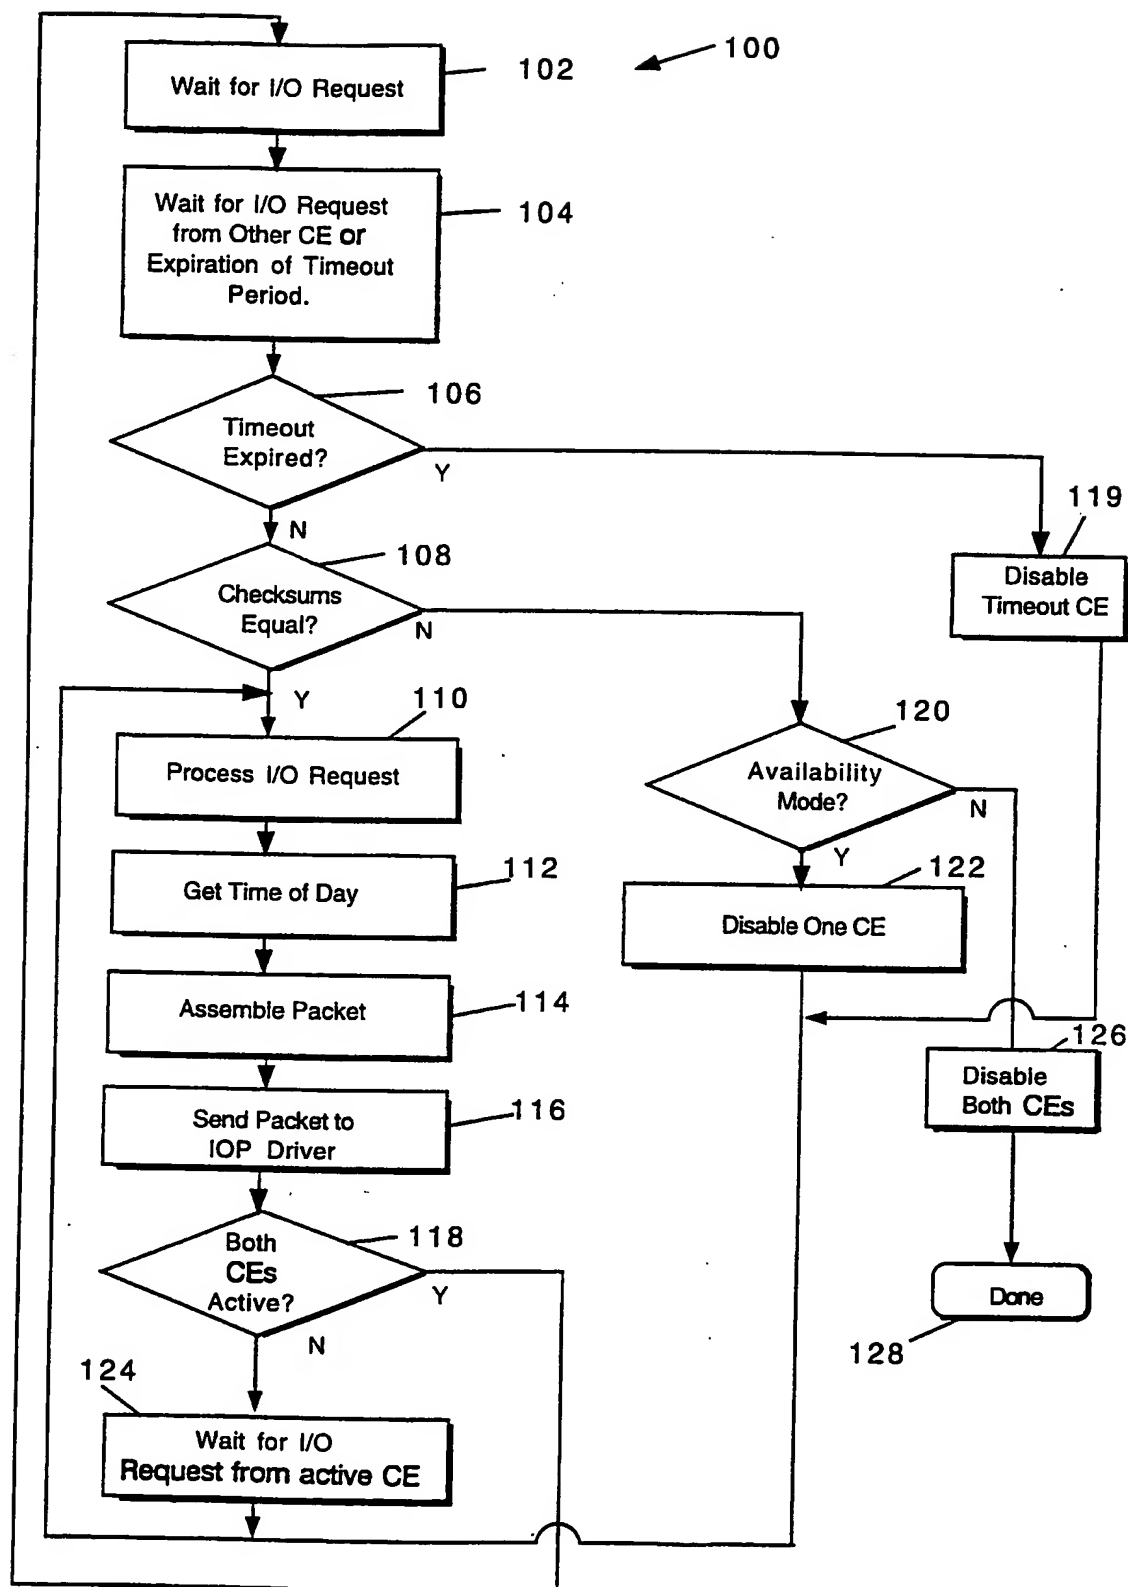


FIG. 3

SUBSTITUTE SHEET (RULE 26)




FIG. 4

SUBSTITUTE SHEET (RULE 26)

5/9




FIG. 5

6/9




FIG. 6




FIG. 7




FIG. 8

8/9

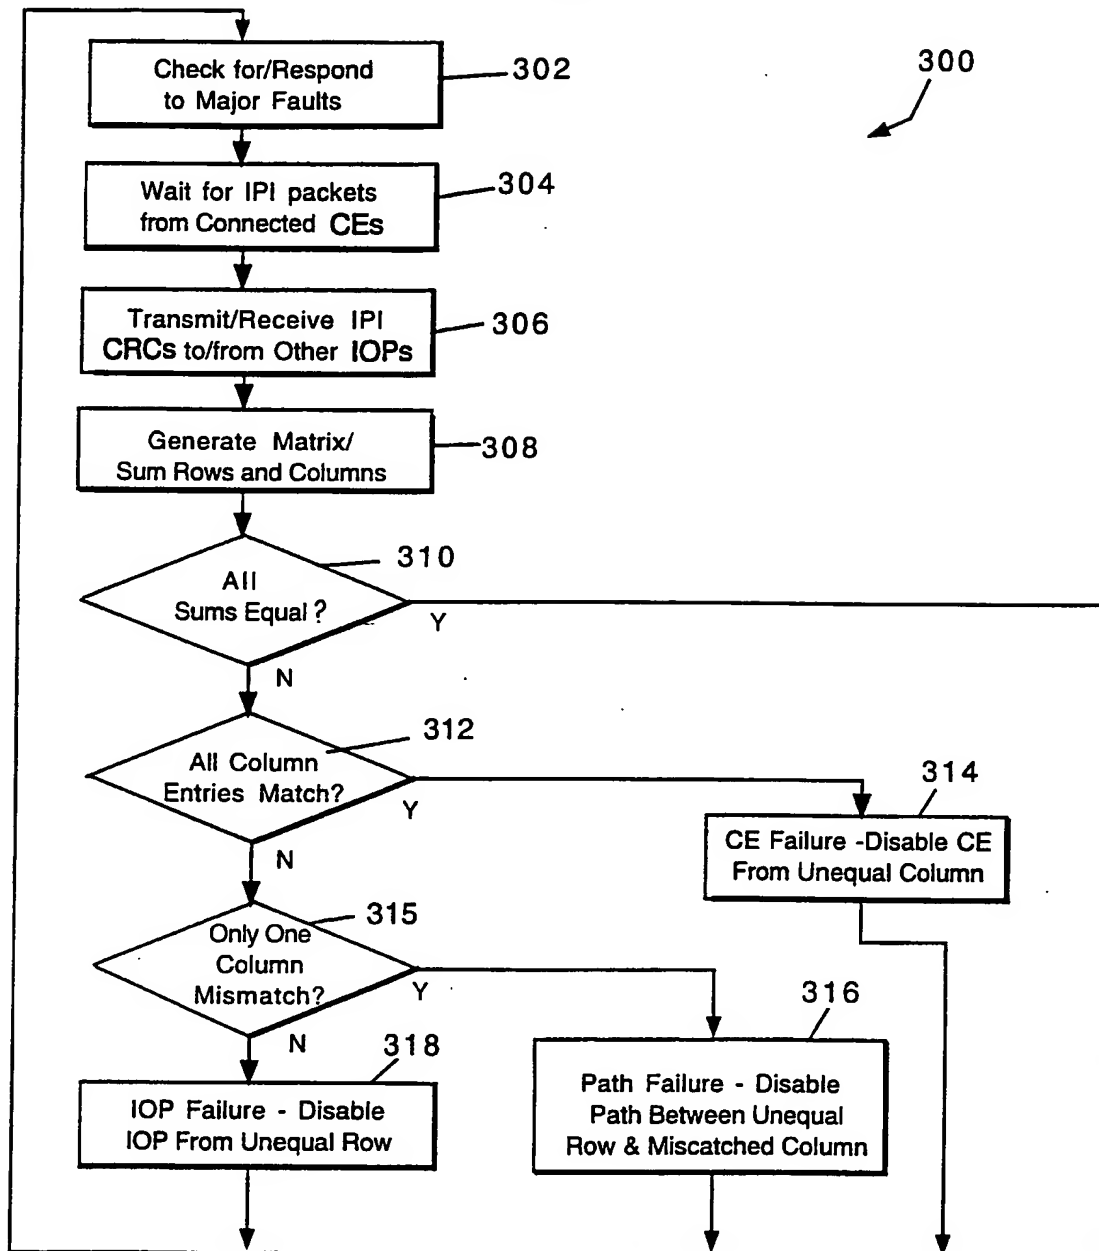


FIG. 9

9/9




FIG. 10

SUBSTITUTE SHEET (RULE 26)

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US94/13350

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 15/16, 11/00, 13/00, 15/76

US CL : 395/550, 575, 275, 800

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/550, 575, 275, 800; 364/270.5, 270.6, 270.7, 371/11.3, 68.3, 36, 9.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Please See Extra Sheet.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ---,P Y	US, A, 5,327,553 (JEWETT ET AL) 5 July 1994 See column 4, line 6 to column 5 line 24, and column 15, line 10 to column 18, line 6, and column 13 line 10 to column 15, line 7.	19,40,42,43,4 5,47,49 ----- 1-18,20-30, 41,44,46, 50
Y	IBM Technical Disclosure Bulletin, volumn 31, No. 11, issued April 1989, "Low Cost Device For contention-Free Barrier Synchronization" See entire document	1-18
Y,P	US, A, 5,339,404 (VANDLING III) 16 August 1994, See Abstract and column 3, lines 26-41 and column 8 line 1 to column 10 line 10	1-4,12,16, 18

☒ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	* T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
* A* document defining the general state of the art which is not considered to be part of particular relevance	* X	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
* E* earlier document published on or after the international filing date	* Y	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
* L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	* A*	document member of the same patent family
* O* document referring to an oral disclosure, use, exhibition or other means		
* P* document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

30 JANUARY 1995

Date of mailing of the international search report

03 MAY 1995

 Name and mailing address of the ISA/US
 Commissioner of Patents and Trademarks
 Box PCT
 Washington, D.C. 20231

Facsimile No. (703) 305-3230

 Authorized officer
Joseph Palys
 JOSEPH PALYS

Telephone No. (703) 305-9600

INTERNATIONAL SEARCH REPORT

International application No.

P S94/13350

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US, A, 5,317,726 (HORST) 31 May 1994 See Abstract and column 2, lines 11-44 and column 8 lines 8-47 and column 16, line 55 to column 22 line 32	1-18
Y	US, A, 5,091,847 (HERBERMANN) 25 February 1992, See column 18, line 26 to column 19, line 44 and column 26, line 15 to column 32, line 15	19,22,20, 21
Y	US, A, 4,823,256 (BISHOP ET AL) 18 April 1989, See Abstract and column 1, line 51 to column 2, line 69, and column 7 lines 5-9, and line 25 to column 8 line 32, and column 9, line 21 to column 10, line 10, and column 17, line 39 to column 23, line 26.	19-22,26, 28-30
Y	US, A, 4,270,168 (MURPHY ET AL) 26 May 1981 See Abstract and column 2, lines 1-61	36-39
Y	US, A, 4,622,667 (YOUNT) 11 November 1986 See Abstract and column 31, line 59 to column 4, line 41 and column 7, line 40 to column 8, line 26	36-39
Y	US, A, 5,261,092 (McLAUGHLIN ET AL) 09 November 1993 See Abstract and column 10, line 32 to column 11 line 24	36-39
Y	US, A, 4,358,823 (McDONALD ET AL) 09 November 1982 See Abstract and column 1, line 59 to column 2, line 34 and column 10 line 24 to line 49	36-39
Y	US, A, 4,920,481 (BINKLEY ET AL) 24 April 1990 See column 5 lines 33-45 and 65-67, and column 6 lines 57-61 and 61-65, and column 7 lines 29-33 and column 9 line 25	31-35
Y	US, A, 4,695,945 (IRWIN) 22 September 1987 See column 2, lines 39-53, and column 6, line 11, and column 9, lines 64-68, and column 10, lines 1-6 and 44-45	31-35
Y	US, A, 4,933,940 (WALTER ET AL) 12 June 1990 See Figure 1 and column 2, line 21 to column 3, line 46 and column 6, line 43 to column 9, line 23 and column 8, line 35 to column 11, line 4 and column 70 lines 10-69	1,2,19,44, 47-50
Y	US, A, 5,166,674 (BAUM ET AL) 24 November 1992, See Abstract and Figure 6, and column 2, line 35 to column 3, line 61, and Column 5, lines 8-38	47-50

INTERNATIONAL SEARCH REPORT

International application No.

B01JUS94/13350

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 4,634,110 (JULICH ET AL) 06 January 1987, See Figures 1,2 and column 2, line 1 to column 4, line 2, and column 4 line 52 to column 7, line 39	40, 47-50
Y	US, A, 4,356,546 (WHITESIDE ET AL) 26 October 1982, See Figures 1,2 and column 2, line 10 to column 3, line 27 and column 5, line 25 to column 8, line 40	40-50

INTERNATIONAL SEARCH REPORT

International application No.

P 94/13350

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Telephone Practice

Please See Extra Sheet.

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐

The additional search fees were accompanied by the applicant's protest.

☐

No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US94/13350

B. FIELDS SEARCHED

Electronic data bases consulted (Name of data base and where practicable terms used):

APS

search terms: time updates, quantum interrupts, asynchronous, synchronous, fault tolerant, suspend or stop operations, resume operations, barrier synchronization.

BOX II. OBSERVATIONS WHERE UNITY OF INVENTION WAS LACKING

This ISA found multiple inventions as follows:

Group I. Claims 1-18, are drawn to controlling processing elements via timing signals, classified in Class 395 Subclass 550.

Group II. Claims 19-30 and 36-39 are drawn to fault detection and recovery from faults in a processing system, classified in Class 395 Subclass 575.

Group III. Claims 31-35 are drawn to the control and operation of I/O components, classified in Class 395 Subclass 275.

Group IV. Claims 40-50, are drawn to the generic combination and/or interconnections of processing components, classified in Class 395 Subclass 800.

The four groups describe different techniques and apparatuses which do not share the same technical features. In particular, Group I, is limited to a method of controlling or regulating a combination of processing elements according to a periodic sequence of timing pulses, Group II is limited to a method of detecting faults in redundant elements, including reconfiguration and/or recovery of the system with regard to a faulty processing element, Group III is drawn to the control and operation of I/O components, and Group IV is drawn to the generic combination and interconnections of processing elements and their peripheral components. All four thus describe different methods as related to special technical features.